Setting Military Reenlistment Bonuses

Dean D. DeWolfe, James G. Stevens, and R. Kevin Wood

Department of Operations Research, Naval Postgraduate School, Monterey, CA 93943

The United States military frequently has difficulty retaining enlisted personnel beyond their initial enlistment. A bonus program within each service, called a Selective Reenlistment Bonus (SRB) program, seeds to enhance reenlistments and thus reduce personnel shortages in critical military occupational specialties (MOSs). The amount of bonus is set by assigning "SRB multipliers" to each MOS. We develop a nonlinear integer program to select multipliers which minimize a function of deviations from desired reenlistment targets. A Lagrangian relaxation of a linearized version of the integer program is used to obtain lower bounds and feasible solutions. The best feasible solution, discovered in a coordinate search of the Lagrangian function, is heuristically improved by apportioning unexpended funds. For large problems, a heuristic variable reduction is employed to speed model solution. U.S. Army data and requirements for FY87 yield a 0-1 integer program with 12,992 binary variables and 273 constraints, which is solved within 0.00002% of optimality on an IBM 3033AP in less than 1.7 seconds. More general models with up to 463,000 binary variables are solved, on average, to within 0.009% of optimality in less than 1.8 minutes. The U.S. Marine Corps has used a simpler version of this model since 1986.

Keywords: Reenlistment bonus, nonlinear integer program, Lagrangian relaxation

Setting Military Reenlistment Bonuses

Dean D. DeWolfe, James G. Stevens, and R. Kevin Wood
Department of Operations Research, Naval Postgraduate School, Monterey, California 93943

The United States military frequently has difficulty retaining enlisted personnel beyond their initial enlistment. A bonus program within each service, called a Selective Reenlistment Bonus (SRB) program, seeks to enhance reenlistments and thus reduce personnel shortages in critical military occupational specialties (MOSs). The amount of bonus is set by assigning "SRB multipliers" to each MOS. We develop a nonlinear integer program to select multipliers which minimize a function of deviations from desired reenlistment targets. A Lagrangian relaxation of a linearized version of the integer program is used to obtain lower bounds and feasible solutions. The best feasible solution, discovered in a coordinate search of the Lagrangian function, is heuristically improved by apportioning unexpended funds. For large problems a heuristic variable reduction is employed to speed model solution. U.S. Army data and requirements for FY87 yield a 0–1 integer program with 12,992 binary variables and 273 constraints, which is solved within 0.00002% of optimality on an IBM 3033AP in less than 1.7 seconds. More general models with up to 463,000 binary variables are solved, on average, to within 0.009% of optimality in less than 1.8 minutes. The U.S. Marine Corps has used a simpler version of this model since 1986. © 1993 John Wiley & Sons, Inc.

The United States' military services have utilized Selective Reenlistment Bonus (SRB) programs since the early 1960s to improve retention of enlisted personnel in specially designated military occupational specialties (MOSs). Examples of MOSs in the U.S. Army include Fighting Vehicle Infantryman, Heavy Antiarmor Weapons Infantryman, Electronic Warfare/Intercept Aviation System Repairer, Avionics Mechanic, etc. The SRB programs are selective in that bonus levels can be set separately for each of several years-of-service intervals within each MOS. The years-of-service intervals are "zones," and we refer to an MOS/zone combination as a "cell." The SRB programs, the framework for which was established by Congress in 1974, are major personnel management tools that encourage eligible personnel to reenlist in their cells instead of choosing to leave military service. The cells in which reenlisting personnel receive bonuses as well as the amount of bonus is determined by the use of "SRB multipliers" of the basic reenlistment bonus. The basic bonus is monthly base pay times number of years of reenlistment which can, within limits, be chosen by the person reenlisting.

Within each MOS, an SRB multiplier can be applied at reenlistment points that fall between 21 months and 14 years of active duty service (ADS). The period between 21 months and 6 years of ADS is designated as Zone A, between 6 and 10 years as Zone B and between 10 and 14 years as Zone C. Personnel ineligible for an SRB are designated for modeling purposes as Zone D.

Each cell has a desired manning level determined by the number of positions required within the organization. Given current manning levels and projected losses, the desired number of cell reenlistments, i.e., a reenlistments target, is established for each cell. An MOS also has a reenlistment target that is just the sum of the targets of its subordinate cells. In many cells and parent MOSs, targets will be met or exceeded without any added inducement. In others, however, an unacceptable target shortfall may exist which can only be reduced by offering a reenlistment bonus. Those cells and MOSs with unacceptable target shortfalls compete for the available budget to acquire nonzero SRB multipliers and thus increase their attractiveness to soldiers eligible to reenlist (assigning a multiplier of 0 is equivalent to no bonus).

Upon reenlistment in a cell designated to receive an SRB, Congressional directives, DoD Instructions 1304.21 (1985) and 1304.22 (1983), authorize to an individual a bonus amount equal to monthly base pay multiplied by the number of years of additional obligated service and by the cell's SRB multiplier, subject to a maximum of $30,000. The service member immediately receives 50% of the bonus with the residual apportioned in equal annual payments over the remainder of the reenlistment. In addition, Congressional directives stipulate that a cell's multiplier may be any value from 0 to 6 and that only 10% of all bonuses awarded may be for amounts greater than $20,000. Each service's present implementation of this policy is more restrictive than the Congressional directives. For example, the U.S. Army allows multipliers in multiples of 0.5 and restricts multipliers to a maximum of 3, 3, and 0 in zones A, B, and C respectively, with a bonus ceiling of $20,000. The U.S. Marine Corps, on the other hand, uses only integer multipliers with a maximum of 5, 4, and 3 in the three zones, respectively, and a bonus ceiling of $16,000.

The overall SRB program objective during each cycle (one year) is, in conjunction with other programs, to minimize the effect of critical shortages on military force readiness. The military pursues this objective by frequently updating the SRB multipliers during each cycle in response to changing needs and resources. During a program update, and prior to this modeling effort, no single objective function existed to evaluate alternate sets of proposed multipliers for the approximately 1000 cells typically eligible for bonuses in the U.S. Army or U.S. Marine Corps. Instead, the office responsible for the SRB program relied primarily on experience to manually assign and judge a particular set of multipliers. For each cycle, a combination of SRB multipliers was subjectively and iteratively modified in a spreadsheet model until it was estimated to satisfy budgetary and Congressional constraints.

The intent of this modeling effort is to provide a much-needed automated decision support tool for the immediate assignment of SRB multipliers, and also to provide an ability for the analysis of real and "what-if" changes to the military force structure and environment that affect an SRB program. The objective of the model is to minimize a function of deviations from reenlistment targets subject to budgetary and Congressional constraints. One approach to this prob-

lem is to model the assignment of SRB multipliers as a nonlinear stochastic mathematical programming problem. Although this approach is of interest and is currently under investigation, it is not without serious difficulties that include the SRB problem's inherent nonlinearity and the large number of decision variables. Therefore, we have chosen to begin by modeling the assignment of SRB multipliers as an integer nonlinear program (INLP) using the estimates of parameters, such as response rates to various multiplier levels, presently available for the manual assignment of SRB multipliers. This essentially deterministic approach is reasonable in light of the environment in which the assignment method is to be used: Close monitoring of program responses, needs and resources allows almost real-time update of parameter estimates and resource data. New SRB multipliers, determined with updated data, can be established multiple times during a cycle. Also, once resources for an SRB program are exhausted or reenlistment objectives are obtained, the program is simply halted and resumed at the beginning of the next cycle.

Recent work by Lovell and Morey [11] on setting monetary inducement levels by MOS for recruitment, as opposed to reenlistment, might appear to be similar to our work. However, their problem and approach are significantly different. They wish to minimize the cost of meeting recruitment targets by allocating enlistment bonuses and college tuition funds to each MOS. The levels of the monetary awards are not decision variables as in the SRB problem, but rather, the decision variables are the number of awards to make for each MOS. Furthermore, the recruitment model has no explicit budget limitation and does not makes use of a response rate function. On the other hand, the use of our methodology to a recruitment problem might be difficult in that our objective function is computationally tractable because it is separable by MOS. Separability is reasonable because the bonus levels in one MOS have little effect on reenlistments in other MOSs as there are few opportunities to switch between MOSs. For the recruitment problem, an objective function analogous to ours would probably not be separable by MOS since setting a large award for one MOS could draw recruits away from enlisting in other MOSs.

To facilitate solution of our basic INLP, it is reformulated as an integer linear program (ILP). A coordinate search of a Lagrangian function, created from the ILP, provides lower bounds and feasible solutions. The best feasible solution, improved with a greedy heuristic to apportion unexpended funds, provides a near-optimal set of SRB multipliers.

The scope of this modeling effort is limited to a single year (cycle), although multiyear research is being conducted (General Research Corporation [9]); no attempt is made to forecast sets of multipliers for subsequent years based on expected retention resulting from the current year's set of multipliers. The model allocates that part of the budget not obligated for annual payments for bonuses from previous years and does not consider the effects of remaining bonus payments that will have to be apportioned in future years. Other assumptions regarding the model are addressed in the succeeding sections of this article.

1. MODEL FORMULATION AND DESCRIPTION

In this section the SRB multiplier problem is formulated first as an INLP using Congressional and budgetary restrictions. Constraint functions are defined

and described and the rationale behind the selection of the objective function is explained. For explanatory and computational purposes, the problem is then converted to an ILP. The model being described in this section is for the U.S. Army. However, development of a model for another service would be similar.

An important assumption of the model is that there are no interactions among MOSs. Interactions exist when reenlistments in one MOS affect another MOS. If it is determined that significant interactions do exist they can be handled by aggregating the affected MOSs. Also, although movement between MOSs can occur during reenlistment (for example, the Army's BEAR program) these effects are considered negligible. Note that the formulation of the model is quite general in that it can accommodate any alternate objective function as long as the assumption of no MOS interactions is maintained.

The decision variables for this model are the SRB multiplier values for each cell. Although Congressional requirements allow fractional multiplier values, the U.S. Army only allows multiplier increments of 0.5. To simplify notation, we model integer multiples of this allowed fraction.

1.1 Model Development

MODEL P1

Indices:

	$h = 1, 2, 3, 4$	zone A, B, C, or D, respectively.
	$l = 1, 2, \ldots, n$	MOSs.
	$t = 3, 4, 5, 6$	years of reenlistment.

Data:

	A_{hl}, A_l	desired manning level in cell hl and MOS l.
	B	available budget for new SRBs.
	C_{hl}	average cost of training personnel to cell hl.
	D_{hl}	reenlistment target for cell hl.
	$\underline{u}_{hl}, \overline{u}_{hl}$	minimum and maximum permissible multiplier in cell hl.
	N_{hl}	number eligible for reenlistment in cell hl.
	P_{hl}	average monthly base pay of an individual in cell hl.
	$R_{hl}(x_{hl})$	reenlistment response rate in cell hl when multiplier x_{hl} is offered.
	$S_{hlt}(x_{hl})$	fraction eligible in cell hl who reenlist for t years when multiplier x_{hl} is offered.
	W_{hl}	exogenous skill weight for cell hl.
	$\delta_{hl}(x_{hl})$	fraction in cell hl who reenlist for more than \$20,000 when multiplier x_{hl} is offered.

Variables:

	x_{hl}	integer representation of the bonus multiplier level in cell hl.

Functions:

	$c_l(x_{1l}, x_{2l}, x_{3l}, x_{4l})$	total penalty associated with MOS l given the multipliers in its subordinate cells.

$d_l(x_{1l}, x_{2l}, x_{3l}, x_{4l})$ penalty associated with MOS l given the multipliers in its subordinate cells.

$r_{hl}(x_{hl})$ cost for setting the multiplier at x_{hl} in cell hl.

$t_{hl}(x_{hl})$ total number of bonus recipients in cell hl due to x_{hl}.

$p_{hl}(x_{hl})$ total number of bonus recipients receiving over $20,000 in cell hl due to x_{hl}.

Formulation:

$$\text{minimize} \sum_l c_l(x_{1l}, x_{2l}, x_{3l}, x_{4l}),$$

$$\text{subject to} \sum_l \sum_h r_{hl}(x_{hl}) \leq B, \tag{1}$$

$$\sum_l \sum_h (p_{hl}(x_{hl}) - 0.1\, t_{hl}(x_{hl})) \leq 0, \tag{2}$$

$$\underline{u}_{hl} \leq x_{hl} \leq \bar{u}_{hl}, \quad \text{and integer } \forall\, h,\, l.$$

In defining the objective function, the model assumes there are no interactions among MOSs. However, each cell's effect on the objective function is dependent upon the performance of the other cells within the same MOS. Therefore, the objective function is just the sum of each MOS's penalty $c_l(x_{1l}, x_{2l}, x_{3l}, x_{4l})$. We note that $x_{4l} \equiv 0$. However, this variable is included in the formulation to reflect the fact that Zone D reenlistments do affect the objective function. The budget constraint (1) regulates the estimated amount of monetary resources allowed across all cells. Finally, the high-level bonus constraint (2) enforces the Congressional requirement that no more than 10% of all SRB bonuses exceed $20,000. At the present time, bonuses are capped at $20,000. However, we model the SRB multiplier problem with the high-level bonus constraint as future competition with civilian industry in employing specially trained personnel, may require bonuses in excess of $20,000.

The resource function,

$$r_{hl}(x_{hl}) = R_{hl}(x_{hl})N_{hl} \times 0.5x_l P_{hl} \left[\sum_t t S_{hlt}(x_{hl}) \right], \tag{3}$$

represents the estimated number of dollars that must be allocated to offer individuals in cell hl an SRB multiplier of x_{hl}, assuming that base pay does not influence reenlistment duration. This function is the product of the estimated number of reenlistments in cell hl, $R_{hl}(x_{hl})N_{hl}$, and 50% of the estimated bonus due upon reenlistment, $0.5x_{hl}P_{hl}[\sum_{t=3}^6 t S_{hlt}(x_{hl})]$. We note that the forms of the resource function and other functions could be simplified if the data were simplified. However, the data are described exactly as they are available and are reported this way for the sake of historical accuracy.

The total bonus function,

$$t_{hl}(x_{hl}) = \begin{cases} 0, & \text{if } x_{hl} = 0 \\ R_{hl}(x_{hl})N_{hl}, & \text{if } x_{hl} > 0 \end{cases}, \tag{4}$$

represents the number of personnel predicted to reenlist in cell hl given a nonzero multiplier. The function $t_{hl}(x_{hl})$ is necessary to calculate the maximum number of personnel allowed bonuses greater than \$20,000. In general, this function is expected to have decreasing slope as the cell's SRB multiplier increases until for some bonus level (probably above the maximum allowed level) all eligible personnel reenlist.

The high-value bonus function,

$$p_{hl}(x_{hl}) = \delta_{hl}(x_{hl})t_{hl}(x_{hl}), \tag{5}$$

represents the number of personnel in cell hl predicted to reenlist for bonus amounts greater than \$20,000. This function is 0 if no one in cell hl can reenlist for greater than \$20,000 and rises to the total number of predicted reenlistments, $t_{hl}(x_{hl})$, when all reenlistment bonuses in cell hl exceed \$20,000.

There is some leeway in defining the objective (penalty) function for P1 and we claim only that our objective function is sensible and, in practice, has given intuitively appealing solutions. Other objective functions or modifications of our objective function can be readily incorporated into our solution methodology. The following list is the set of guidelines we used in developing our objective function.

(a) If estimated responses for each cell in MOS l exactly meet target values D_{hl} for each cell, the penalty, $c_l(x_{1l}, x_{2l}, x_{3l}, x_{4l})$, should be 0.
(b) As estimated deviations from cell targets rise, so should the penalty. Furthermore, the penalty should rise superlinearly in these deviations so as to spread shortages among cells.
(c) Standard military policy allows some assignment of personnel within an MOS to jobs for which they may be over- or underqualified. This will sometimes correspond to personnel in one cell of an MOS being assigned to a position which ideally should be filled by someone in another cell of that MOS. In some way then, cell deviations in an MOS at full strength should be penalized less than in an MOS which is understrength. Consequently, the penalty function must also reflect deviations from the MOS target defined as $\Sigma_h D_{hl}$.
(d) Some normalization of penalties by desired cell and MOS manning levels must take place since a shortage of 5 people in a cell or MOS containing 1000 people is more easily tolerated than a shortage of 5 in a cell or MOS containing 100. (The importance to the military of a cell or MOS is not particularly dependent on its nominal size.)
(e) The alternative to enticing people to reenlist with a high SRB multiplier is to recruit and train new personnel. Consequently, given two cells of identical size and shortfall, the cell with the higher training cost should have a larger penalty and hence have a greater propensity to be assigned a positive multiplier.

Each MOS's penalty function,

$$c_l(x_{1l}, x_{2l}, x_{3l}, x_{4l}) = \sum_h (|D_{hl} - R_{hl}(x_{hl})N_{hl}|^p)(C_{hl}W_{hl}/A_{hl})(d_l(x_{1l}, x_{2l}, x_{3l}, x_{4l})), \tag{6}$$

where

$$d_l(x_{1l}, x_{2l}, x_{3l}, x_{4l}) = 1 + \left| \sum_h (D_{hl} - R_{hl}(x_{hl})N_{hl}) \right| \bigg/ A_l, \qquad (7)$$

is the sum of a product of three sets of terms. The first term, resulting from guidelines (a) and (b),

$$|D_{hl} - R_{hl}(x_{hl})N_{hl}|^p,$$

represents the deviation from the desired number of reenlistments in each cell hl raised to the pth power. A cell's deviation is calculated by taking the absolute difference between the desired number of reenlistments D_{hl} and the estimated number of reenlistments $R_{hl}(x_{hl})N_{hl}$. For a cell with a shortfall, the cell's deviation is decreasing as reenlistments rise due to increasing the SRB multiplier until, for some multiplier, it may begin to increase as the shortfall becomes a surplus. Any reasonable value for $p > 1$ has the desired effect of penalizing those cells with large deviations much more severely than smaller ones, which tends to even out shortages and surpluses among cells. For our purposes we have chosen $p = 2$, as this yields objective function units of weighted dollars. We also note that it is a simple generalization to allow shortages and surpluses to be weighted differently but we ignore this for notational simplicity.

Given the cell deviation, the second term $C_{hl}W_{hl}/A_{hl}$ incorporates the training cost as referred to in guideline (e), a subjective weight, and normalization by desired cell manning level as per guideline (d). The cost C_{hl} is the term presently used in the manual assignment of SRB multipliers to compare target deviations across cells. In addition, a few cells may be considered more critical than others due to special attributes other than cost, e.g., special operational forces. The weight W_{hl} is an exogenous factor which permits the user to consider special cell attributes and manually override other model factors. However, in most cases we expect W_{hl} to be set to a nominal value of 1 and we do this in all computations reported in this article. Finally, division by A_{hl} normalizes the deviation of a cell with respect to its ideal strength. The product of the first two terms of (6) is then the "basic" penalty function for cell hl.

The third term (7) is a measure of deviation from the desired manning in MOS l. This unitless term modifies the effect of each cell's basic penalty within an MOS and is motivated by guideline (c) and normalized as per guideline (d). If the MOS, taken as a whole, is on target, then the sum of basic cell penalties define the penalty associated with the MOS. However, if the MOS also deviates from desired strength this worsens the individual cell penalties.

The response rate of a cell hl to an SRB multiplier $R_{hl}(x_{hl})$ is critical in predicting each cell's expected deviation and resource requirements. We used the same response rate estimates as the U.S. Army uses in their manual computation of SRB multipliers. More sophisticated methods for determining response rates (Cymrot [2]), which take into consideration the effects of demographics and economic conditions, have been implemented by the Marine Corps and are under study for adoption by the U.S. Army (Streff [12]). These results

should improve the trustworthiness of response rate estimates and are easily incorporated into this model.

Data required for estimating the fraction of reenlistment eligible soldiers in cell hl who reenlist for t years when multiplier x_{hl} is offered, $s_{hlt}(x_{hl})$, are available, but have not yet been compiled; consequently we estimate it from the average duration of reenlistment in cell hl. As indicated, $s_{hlt}(x_{hl})$ is dependent on x_{hl}. If x_{hl} is permitted to increase to levels that allow SRB bonuses greater than \$30,000, we expect the majority of personnel to reenlist for a duration not to exceed the time required to maximize their SRB.

As presently modeled, P1 is an INLP. The special structure of this problem and the nonlinearity in the objective function and both complicating constraints make solution by standard branch-and-bound or dynamic programming impractical or impossible. Given these difficulties, we reformulate the problem into an ILP for solution using a specialized procedure.

1.2. Integer Linear Program Conversion

Here, the model is transformed from an INLP to an ILP to facilitate its solution. The number of explicit constraints does increase by the number of MOSs, but this causes little difficulty with our solution approach. To facilitate notation let the vector $\bar{m}$ be a 4-tuple with range encompassing potential multipliers in the four zones of any MOS, and M_l represent the set of allowable multipliers in MOS l; i.e., $M_l = \{\bar{m} = (m_1, m_2, m_3, m_4): \underline{u}_{hl} \leq m_h \leq \bar{u}_{hl}, m_h$ integer, $h = 1, 2, 3, 4\}$. Then, the ILP equivalent of Model 1 is:

MODEL P2

Indices:

$\bar{m}$ A vector whose values represent a multiplier set in an MOS. $\bar{m} = (m_1, m_2, m_3, m_4)$.

$h = 1, 2, 3, 4$ zone A, B, C, or D, repsectively.

$l = 1, 2, \ldots, n$ MOS.

Functions:

$c_{\bar{m}l}$ $c_l(m_1, m_2, m_3, m_4)$

$r_{\bar{m}l}$ $\sum_h r_{hl}(m_h)$

$t_{\bar{m}l}$ $\sum_h t_{hl}(m_h)$

$p_{\bar{m}l}$ $\sum_h p_{hl}(m_h)$

Variables:

$y_{\bar{m}l}$ A binary variable representing the multiplier levels in the cells of MOS l. For $\bar{m} \in M_l$, $y_{\bar{m}l} = 1$ if the Zone A multiplier is m_1, the Zone B multiplier is m_2, the Zone C multiplier is m_3, and the Zone D multiplier is m_4. Otherwise $y_{\bar{m}l} = 0$.

Formulation:

$$\text{minimize} \quad \sum_l \sum_{\bar{m} \in M_l} c_{\bar{m}l} y_{\bar{m}l},$$

$$\text{subject to} \quad \sum_l \sum_{\bar{m} \in M_l} r_{\bar{m}l} y_{\bar{m}l} \le B, \tag{8}$$

$$\sum_l \sum_{\bar{m} \in M_l} (p_{\bar{m}l} - 0.1\, t_{\bar{m}l}) y_{\bar{m}l} \le 0, \tag{9}$$

$$\sum_{\bar{m} \in M_l} y_{\bar{m}l} = 1, \quad \forall \ l, \tag{10}$$

$$y_{\bar{m}l} \in \{0, 1\}, \quad \forall \ \bar{m} \in M_l, l.$$

2. SOLUTION METHODOLOGY

Theoretically, the SRB model P2 could be solved by a standard, LP-based branch-and-bound algorithm. However, given the potential size of the model (10,000 to 500,000 0-1 variables), and the desire to implement the model on a microcomputer, this approach is not attractive. The SRB model is a generalization of the multi-item scheduling model of Sweeney and Murphy [13], in that their model contains "multiple choice" constraints (10) and a single budget constraint (8) but no additional complicating constraints such as (9). Their branch-and-bound method could be generalized to the SRB model, but their method would also require solving the LP relaxation of P2. Also, Bean [1] describes a similar method for solving problems of the form of P2 but, once again, solving the LP relaxation of P2 would be required. We also note that the computational experiences discussed by Sweeney and Murphy and by Bean are limited to problems with at most a few thousand 0-1 variables.

In order to avoid solving large LPs we have taken an approach to solving P2 based on solving a Lagrangian relaxation of that model (see Fisher [7] for an overview of Lagrangian relaxation). We invariably obtain an excellent lower bound, equivalent to the LP-based lower bound, from this solution technique. In the process we also obtain feasible solutions to P2 which are heuristically improved to yield a high-quality final solution to P2. Very modest optimality gaps have obviated the need for any sort of enumeration approach.

Let $\lambda = (\lambda_1, \lambda_2)$. The relaxed formulation of P2 is

$$\max_{\lambda \ge 0} \left\{ \begin{array}{l} \min_y \ \sum_l c_{\bar{m}l} y_{\bar{m}l} + \lambda_1 \overbrace{\left(\sum_l r_{\bar{m}l} y_{\bar{m}l} - B \right)}^{1^{**}} + \lambda_2 \overbrace{\left(\sum_l (p_{\bar{m}l} - 0.1\, t_{\bar{m}l}) y_{\bar{m}l} \right)}^{2^{**}} \\ \text{s.t.} \quad \sum_{\bar{m} \in M_l} y_{\bar{m}l} = 1, \quad \forall \ l \\ \qquad\quad y_{\bar{m}l} \in \{0, 1\}, \quad \forall \ \bar{m} \in M_l, l \end{array} \right\}$$

$$\text{(LR-P2)}$$

where the inner portion of the objective function may be equivalently written

$$F(\lambda) = \min_{y} \sum_{l} (c_{\overline{m}l} + \lambda_1 r_{\overline{m}l} + \lambda_2(p_{\overline{m}l} - 0.1 \ t_{\overline{m}l})) y_{\overline{m}l} - \lambda_1 B. \tag{11}$$

For fixed values of λ_1 and λ_2, the inner minimization (11) selects the set of multipliers that corresponds to the subproblem

$$\min_{\overline{m} \in M_l} \{c_{\overline{m}l} + \lambda_1 r_{\overline{m}l} + \lambda_2(p_{\overline{m}l} - 0.1 \ t_{\overline{m}l})\}, \tag{12}$$

for each MOS l. The worst subproblem in the case of the U.S. Army would have $|M_l| = 7^2 = 49$ different values of $\overline{m}$ to investigate corresponding to an MOS with allowed multiples of 0.5 from 0 through 3 for zones A and B with multipliers fixed at 0 for zones C and D. Note here also that the optimal solution value for LR-P2, denoted $F(\lambda^*)$, is equivalent to the lower bound from the LP relaxation of problem P2 since the LP solution to the inner minimization of LR-P2 is intrinsically integer (Geoffrion [10]).

The outer maximization problem LR-P2 is concave and could be solved using subgradient optimization; however, obtaining convergence with subgradient optimization can be difficult in practice. Therefore, a coordinate search of the Lagrangian function is implemented for the two dimensions of this model. In Section 2.1 we briefly discuss the coordinate search procedure and point out some special techniques used to make this procedure efficient. Since $F(\lambda^*)$ obtained from the coordinate search is equivalent to the LP lower bound for P2, $F(\lambda^*)$ might be a poor bound and we might be unable to prove that the feasible solution we obtain is good because of a poor bound. However, in Section 2.2 we provide theoretical evidence that, because of the SRB problem's special structure, the bound should be good: We show that the ratio of the LP lower bound to the value of the optimal integer solution converges to 1 as the number of MOSs becomes large. Furthermore, although we cannot show that our heuristic will always give a good solution, we can show that an "LP rounding" heuristic will give a good solution which lends credibility to the use of a heuristic approach.

2.1. Coordinate Search Procedure

The lower bound for λ is zero. When $\lambda = 0$ the multiplier set chosen for each MOS corresponds to the set of multipliers whose attributes have minimum penalty. If P2 is feasible when $\lambda = 0$ (a highly unlikely situation in practice), then $F(0) = F(\lambda^*)$ and we are done. If $\lambda = 0$ fails to provide a feasible solution for P2, then for values of λ_1 and λ_2 sufficiently large, the solution to (12) for each MOS is a multiplier set with attributes having minimum cost, minimum bonus constraint contribution, and maximum penalty; i.e., the model would choose MOS multiplier sets that correspond to the minimum permissible SRB levels for all cells. If this solution is feasible, an initial upper bound for P2 has been established and a coordinate search is performed to identify λ^*. Values for λ_1 and λ_2 great enough to force the selection of minimum cell multipliers for all

cells without obtaining a feasible solution for P2 indicate that a $\lambda \geq 0$ does not exist that provides a feasible solution and the problem is trivially infeasible. This situation will only arise if the minimum cell multipliers require a violation in either the budget or bonus constraint. However, in practice this does not occur since, usually $\underline{u}_{hl} = 0$ for all h and l.

The coordinate search on the Lagrangian function repetitively fixes one of the two components of λ while finding the value of the other component that maximizses $F(\lambda)$. Once each component is bounded, the optimization in each coordinate can be carried out by bisection search given the slope of $F(\lambda)$ in the appropriate coordinate. These slopes are simply the values of 1^{**} for the first coordinate and 2^{**} for the second coordinate. Because the one-dimensional optimizations are so simple, it is also possible to use a very simple cutting plane algorithm in place of the bisection search. The "master problem" of the cutting plane algorithm consists of calculating the intersection of two lines. Empirically, we have found that the most efficient approach is to use the cutting plane algorithm for the first few iterations and then switch to the bisection search.

Throughout the coordinate search, those feasible sets of SRB multipliers encountered while solving for λ^* provide upper bounds for P2. The current best known set of multipliers is stored as the incumbent solution to P2 with an attempt made to improve the last such solution using a marginal rate-of-return heuristic as described in Section 4.

2.2. Model P2's Special Structure

The special structure of P2 consists of the generalized upper bound (GUB) constraints (10) and the fact that if the LP relaxation of P2 is feasible, a feasible integer solution to P2 can always be obtained. That this is true can be seen as follows. Suppose we are given an extreme point solution to the LP-relaxation of P2 and that for MOS l, $y_{\overline{m}l}$ is fractionated for $\overline{m} = \overline{m}^1, \overline{m}^2, \ldots, \overline{m}^j$ (for our problem j is at most 3). Let $\overline{m}' = (m_1', m_2', m_3', m_4')$ where $m_h' = \min\{m_h^1, m_h^2, \ldots, m_h^j\}$. Now, $\overline{m}' \in M_l$ and setting $y_{\overline{m}'l} = 1$ for MOS l will always contribute no more to the left-hand sides of (8) and (9) than will the original fractionated solution and is thus feasible. This process could be repeated for each MOS with fractionated variables and will be referred to as the "rounded LP solution".

Let G_{nk}^{IP} and G_{nk}^{LP} represent the optimal objective function values for P2 and its LP relaxation, respectively, for an SRB problem with n MOSs and k complicating constraints. Because of the GUB structure of constraint (10), k is also the maximum number of such constraints in which fractionation can occur in the optimal solution to the LP relaxation of P2. In addition, let G_{nk}^{LP-} be that portion of G_{nk}^{LP} corresponding to the $n - k$ or more variables that are integer valued in the optimal solution. By construction, and since P2 is a minimization problem, $0 \leq G_{nk}^{LP-} \leq G_{nk}^{LP} \leq G_{nk}^{IP}$. Next, let g_l represent the difference between the maximum and minimum objective function contributions for allowable multiplier sets within MOS l. This value bounds the impact to the objective function of using the rounded LP solution over the continuous LP solution for MOS l. We assume that g_l is finite and let $\overline{g}$ represent the supremum over l of g_l.

PROPOSITION 1: If as $n \to \infty$, $G_{nk}^{LP} \to \infty$, then $G_{nk}^{IP}/G_{nk}^{LP} \to 1$.

PROOF: Since $\bar{g}$ is nonnegative, finite, and bounds the worst possible objective function increase for each fractionated MOS "rounded" to integrality then, $G_{nk}^{LP-} \le G_{nk}^{LP} \le G_{nk}^{IP} \le G_{nk}^{LP-} + k\bar{g}$ which implies $G_{nk}^{LP-} \to \infty$ as $n \to \infty$. But this also implies that as $n \to \infty$,

$$1 \le \frac{G_{nk}^{IP}}{G_{nk}^{LP}} \le \frac{G_{nk}^{LP-} + k\bar{g}}{G_{nk}^{LP}} \le \frac{G_{nk}^{LP-} + k\bar{g}}{G_{nk}^{LP-}} = 1 + \frac{k\bar{g}}{G_{nk}^{LP-}} \to 1.$$

Q.E.D.

Therefore, the LP and IP solutions for P2 will be close for large n.

From the above discussion and proof it is clear that, for large n, a good approximate solution to P2 could be obtained using the "LP rounding procedure". Although we will not pursue this procedure, since it requires solving the LP relaxation, it is possible to obtain an a priori (before solving the model but after establishing its coefficients) bound on the error of such a solution. This bound may be used as a performance measure of our Lagrangian relaxation procedure. Let G_{nk} be the objective value of the rounded LP solution of P2. Also, let $\underline{G}_{nk}^{LP}$ represent P2's optimal objective function value without complicating constraints (8) and (9). Define the error of the LP rounding solution as

$$\text{err} = \frac{G_{nk} - G_{nk}^{IP}}{G_{nk}^{IP}}.$$

PROPOSITION 2: An a priori upper bound for error for the LP rounding solution of P2 is

$$\text{err} \le \frac{k\bar{g}}{\underline{G}_{nk}^{LP}}.$$

PROOF: Since $G_{nk}^{LP} \le G_{nk}^{LP} \le G_{nk}^{IP} \le G_{nk} \le G_{nk}^{LP-} + k\bar{g} \le G_{nk}^{LP} + k\bar{g}$,

$$\text{err} = \frac{G_{nk} - G_{nk}^{IP}}{G_{nk}^{IP}} = \frac{G_{nk}}{G_{nk}^{IP}} - 1 \le \frac{G_{nk}^{LP} + k\bar{g}}{G_{nk}^{IP}} - 1 \le \frac{G_{nk}^{LP} + k\bar{g}}{G_{nk}^{LP}} - 1 = \frac{k\bar{g}}{G_{nk}^{LP}} \le \frac{k\bar{g}}{\underline{G}_{nk}^{LP}}.$$

Q.E.D.

For the SRB problem with U.S. Army FY87 data and Congressional restrictions, $\bar{g}$ is approximately 0.08% of $\underline{G}_{nk}^{LP}$ for $n = 272$. This yields an error bound of approximately 2% for the LP rounding procedure.

3. VARIABLE REDUCTION

The model may operate in two different environments depending upon the maximum bonus multiplier and maximum monetary bonus parameters. A level of either parameter that restricts the maximum monetary bonus to $20,000 or less effectively eliminates the high-level bonus constraint from the model. If this

is the case, LR-P2 is reduced to

$$\max_{\lambda_1 \geq 0} \left\{ \begin{array}{l} \min_{y} \quad \sum_{l} \sum_{\overline{m} \in M_l} c_{\overline{m}l} y_{\overline{m}l} + \lambda_1 \left(\sum_{l} \sum_{\overline{m} \in M_l} r_{\overline{m}l} y_{\overline{m}l} - B \right) \\ \text{s.t.} \quad \sum_{\overline{m} \in M_l} y_{\overline{m}l} = 1, \quad \forall \ l \\ \qquad\qquad y_{\overline{m}l} \in \{0, 1\}, \quad \forall \ \overline{m} \in M_l, l \end{array} \right\} \qquad (\text{LR-P2R})$$

3.1. Reduction of Case 1—LR-P2R

Analysis of the objective function indicates that variables can exist in LR-P2R that are dominated, i.e., will have a value of 0 in the optimal solution. Elimination of these dominated variables reduces model storage requirements and can save time if the time for the variable reduction (e.g., Garfinkel and Nemhauser [8]) plus the reduced model's execution time is less than the execution time of the full model (variable reduction in a typical U.S. Army problem resulted in a 25% reduction in solution time and a 75% reduction in storage).

Each MOS l in LR-P2R contains a variable $\underline{y}$ representing the minimum bonus multiplier levels permitted in its subordinate cells. This variable has penalty $\underline{c}$ and the minimum cost $\underline{r}$ of all variables associated with MOS l. Any variable in MOS l with greater penalty than $\underline{c}$ cannot be in LR-P2R's optimal solution because its cost is necessarily greater than $\underline{r}$. A second reduction in the number of variables can be obtained by observing that increasing cost must be offset by decreasing penalty. Since this variable reduction does not affect the optimal solution $F_r(\lambda_r^*) = F_r(\lambda^*) = F(\lambda^*)$, where F_r represents the objective function in LR-P2R and λ_r^* represents the optimal λ for LR-P2R given the reduced variable set. Again this only occurs when monetary bonuses are restricted to $20,000 or less.

3.2. Reduction of Case 2—LR-P2

If multiplier levels and bonus ceilings allow bonuses over $20,000, only those variables with less penalty and a greater percentge of reenlistments with bonuses over $20,000 than $\underline{y}$ are dominated. Typically, however, this reduction is small and the manipulation of the resulting variable set to solve LR-P2 (up to 463,000 variables for a single perturbation of the Army's FY87 data) suggested a heuristic to reduce the number of model variables and solution time while finding a solution near $F(\lambda^*)$. Of the several schemes tested the one resulting in considerable time improvement and whose solution approached that of the unrestricted model was to implement the variable reduction outlined in the previous section, i.e., restrict initial model solution to those variables in each MOS whose penalty was less than $\underline{c}$ and whose increased cost was offset by decreased penalty. Let λ_r^* be the optimal λ for LR-P2 using the heuristically reduced set of variables and let $F_r(\lambda_r^*)$ represent the corresponding objective function value. Then if $\lambda_2 = 0$ at λ_r^*, $F_r(\lambda_r^*) = F(\lambda^*)$, and we are done; i.e., the heuristic variable reduction did not eliminate any variable in the optimal solution of LR-P2. How-

ever, if $\lambda_2 \neq 0$ at λ_r^*, then $F(\lambda_r^*) \leq F_r(\lambda_r^*)$ and the user has an option to accept the restricted solution $F(\lambda_r^*)$ or use λ_r^* as an initial guess to continue the search for $F(\lambda^*)$.

4. MARGINAL RATE-OF-RETURN HEURISTIC

It seems reasonable to expect that the best feasible solution obtained in the process of optimizing LR-P2 would require the expenditure of nearly the entire SRB budget. There might, however, be some budget dollars remaining that could be apportioned to further reduce the objective value of the best feasible solution (the present upper bound) if the model constraints simultaneously remain feasible. Accordingly, a heuristic method for allocating residual budget resources was developed using the concept of marginal rate of return (Everett [6]).

Having obtained the optimal set of SRB multipliers from LR-P2*, the additional cost due to incrementing each MOS's multiplier set to its next best penalty, i.e., increase each MOS's budget allocation and thus decrease its penalty, is calculated for all those MOSs with multiplier sets not already at their maximum values. If the incremental amount required for an MOS is less than the unallocated portion of the budget and would not cause a violation in the high-level bonus constraint, it is possible to increment that MOS's multiplier set and remain feasible. Thus, the incremental improvement in the objective function per dollar spent in MOS l is defined to be

$$\mathrm{ROR}_l = \frac{c_{\overline{m}l} - c'_{\overline{m}l}}{r'_{\overline{m}l} - r_{\overline{m}l}},$$

where $c_{\overline{m}l}$, $r_{\overline{m}l}$ and $c'_{\overline{m}l}, r'_{\overline{m}l}$ represent the penalty and resource requirements associated with the best yet and next best multiplier set for MOS l, respectively. In the heuristic procedure, the MOS with the maximum ROR_l that does not cause a violation in the model constraints has its multiplier set modified. The budget and high-level bonus constraints are updated and the process is repeated until it is no longer possible to decrease the penalty of any MOS and remain feasible. The resulting set of SRB multipliers gives an upper bound for the SRB multiplier problem and is accepted as the final solution.

One problem can arise with the above procedure: In optimizing the Lagrangian function, no feasible solution may ever be found. In fact, this occurred in 1 of 50 test problems described in the next section. In this case we have implemented a "deletion heuristic" to find a good solution. This heuristic uses the concept of marginal rate of return in the reverse of what is described above, i.e., starting with the infeasible solution at λ^*, the multiplier set is reduced for that MOS which yields the greatest decrease in infeasibility for the least increase in penalty. This is repeated until feasibility is obtained.

Other heuristic improvement procedures are possible but were not implemented in this study. For example, by simultaneously incrementing one MOS's multiplier set while decrementing another's, it might be possible to improve the upper bound established with the marginal rate-of-return heuristic while also remaining feasible with respect to both the bonus and high-level budget con-

straints. In fact, problem P2 could be solved to optimality using a branch-and-bound algorithm, though the large number of decision variables might make this intractable. Pursuing strict optimality with the resultant increase in algorithmic complexity and solution times was not deemed worthwhile in light of the good solutions that have been consistently obtained for various data sets and objective functions.

5. COMPUTATIONAL RESULTS

U.S. Army requirements for FY87 (bonus ceiling of $20,000 and maximum SRB multipliers of 3, 3, and 0 in zones A, B, and C respectively) effectively eliminate the high-level bonus constraint (9) in P2 and yield a model with 12,992 binary variables and 273 constraints. This model is solved in 1.67 seconds on an IBM 3033AP under the VM/CMS operating system using VS Fortran 77. Model construction, including data manipulation, variable reduction, and other functional computations, required 1.50 seconds. The resulting model of 3072 variables was solved to within 0.0005% error in 0.15 seconds and required 99.97% of the authorized budget. In 0.02 seconds the marginal rate-of-return heuristic improved the solution to within 0.00002% and utilized 99.99% of the budget.

Because the solution procedure is a heuristic and might be sensitive to data changes, FY87 data obtained from the U.S. Army Personnel Center were perturbed to determine model robustness. Test 1 consisted of 100 model runs and restricted model execution to the U.S. Army requirements described above. Again the U.S. Army restrictions effectively eliminate the high-level bonus constraint (9) and produce 0-1 integer programs with 273 constraints and 12,992 variables. The data: A_{hl}, D_{hl}, E_{hl}, P_{hl}, $R_{hl}(x_{hl})$, and C_{hl} were perturbed as follows:

$$d^* = d \times U, \tag{13}$$

where d^* is perturbed datum,
  d is original datum, and,
  U is uniform (0.5, 1.5) random variate.

Any data originally integer, such as reenlistment targets, were rounded down to the nearest integer after being perturbed. In addition, reenlistment response rates were necessarily capped at 1. Model error was computed as follows:

$$e = \frac{\overline{F} - \underline{F}}{\underline{F}}, \tag{14}$$

where $\overline{F}$ and $\underline{F}$ represent, respectively, the best feasible solution and the maximum Lagrangian function evaluation of LR-P2 discovered during model solution.

Test 1 solutions with variable reduction required an average of 1.68 seconds for each run with a maximum run of 1.84 seconds. Without variable reduction the average time of solution and storage requirements increased 25% and 400%, respectively. Approximately 90% of the solution time is utilized for data manipulation, variable reduction, and the functional computations required to construct the model (denoted "Development" in Table 1). Given the reduced vari-

Table 1.   Summary of the results for U.S. Army FY87 data and restrictions perturbed 100 times.

Run	Time (sec)				Budget		Error	
	Development	Solver	Heuristic	Total	Solver	Final	Solver	Final
Average	1.50	0.16	0.02	1.68	99.40%	99.99%	0.00878%	0.00025%
Best	1.38	0.11	0.00	1.52	99.99%	99.99%	0.00012%	0.00000%
Worst	1.56	0.30	0.06	1.84	97.12%	99.99%	0.04830%	0.00274%

able set, an average of 0.16 seconds was needed to discover λ^*, the optimal Lagrangian solution (denoted "Solver"). In a negligible amount of time, the marginal rate-of-return heuristic (denoted "Heuristic") improved the average initial solution's error from 0.0088% to 0.0003% and increased budget expenditure from 99.40% to 99.99%. Other Test 1 results are included in Table 1.

In order to test the full generality of the solution procedure, a relaxed version of the U.S. Army problem was solved. Test 2 permitted model execution up to Congressional restrictions, i.e., a maximum bonus of $30,000 with allowable SRB multipliers of 6, 6, and 6 in zones A, B, and C. Thus, both constraints (8) and (9) in P2 were enforced during the solution procedure. Test 2 consisted of 50 model runs and yielded 0-1 integer programs with 274 constraints and up to 463,000 variables. In Test 2, model solutions with heuristic variable reduction required an average of 106.22 seconds with a maximum run of 166.72 seconds (solution times without heuristic variable reduction ranged from 11 minutes to 2 hours using an in-core/out-of-core algorithm). Approximately 70% of the solution time was utilized for data manipulation, heuristic variable reduction, functional computations, and storage of the variable set for the evaluation of $F(\lambda_r^*)$ ("Development"). The discovery of the Lagrangian solution $F_r(\lambda_r^*)$ and the evaluation of $F(\lambda_r^*)$ required an average of 28.66 seconds (Solver). The majority of this time was for evaluating $F(\lambda_r^*)$ using information stored out of core. Only one of 50 runs required a deletion heuristic (elimination of costly SRB multiplier sets) to be used on the solution at λ_r^* to improve the run's upper bound. For all runs the error even prior to the marginal rate-of-return heuristic was less than 0.14%, which did not warrant a continued search for λ^*. Again, using a negligible amount of time, the marginal rate-of-return heuristic improved the average initial solution's error from 0.0304% to 0.0088%. Other Test 2 results are included in Table 2.

Table 2.   Summary of the results for U.S. Army FY87 data and Congressional restrictions perturbed 50 times.

Run	Time (sec)				Budget		Error	
	Development	Solver	Heuristic	Total	Solver	Final	Solver	Final
Average	77.55	28.66	0.02	106.22	99.02%	99.99%	0.03041%	0.00878%
Best	38.22	0.24	0.01	53.09	99.99%	99.99%	0.00034%	0.00002%
Worst	140.44	58.75	0.05	166.72	93.35%	99.99%	0.13947%	0.09615%

6. CONCLUSIONS

Selective Reenlistment Bonus (SRB) Programs are major personnel management tools that encourage enlisted personnel to reenlist in their critical MOS and zone combinations instead of leaving military service. In this article an objective function to measure a program's effectiveness is developed together with a mathematical program and marginal rate-of-return heuristic that approximately optimize the objective function subject to budgetary and high-level bonus constraints. A preliminary version of this model, implemented on a microcomputer, has been used by the U.S. Marine Corps since 1986 (DeWolfe [3]). The new model is an improvement over the Marine Corps model in that it considers cell manning dependent upon the manning of the parent MOS and also allows consideration of the complicating high-level bonus constraint. The new model is currently being implemented by the U.S. Army and will initially be used for determining future SRB program budgets.

The model and solution methodology are very general and only require that interactions among MOSs be prohibited. Other penalty and resource utilization functions are readily accommodated by this model. Improvements in response rate estimates (Cymrot [2] and Streff [12]), should further increase the accuracy with which the model's multipliers achieve their predicted cell targets.

The future portends increasingly restrictive budgets within the Department of Defense, additional pressure to maximize benefits given limited resources, and intense competition with civilian industry to maintain specially trained personnel within the military. With respect to maximizing the benefits realized from an SRB program, a nearly optimal assignment of SRB multipliers with respect to budgetary and Congressional restrictions goes far towards achieving that goal.

REFERENCES

[1] Bean, J.C., "A Lagrangian Algorithm for the Multiple Choice Integer Program," *Operations Research*, **32**, 1185–1193 (1984).

[2] Cymrot, D.J., "The Effects of Selective Reenlistment Bonuses on Retention," Center for Naval Analysis Research Memorandum No. 87-43, March 1987.

[3] DeWolfe, D.D., "Determination of Selective Reenlistment Bonus Multipliers in the United States Marine Corps," M.S. thesis, Naval Postgraduate School, March 1986.

[4] DoD Instruction 1304.21, "Award of Enlisted Personnel Bonuses and Special Duty Assignment Pay," August 2, 1985.

[5] DoD Instruction 1304.22, "Award of Enlisted Personnel Bonus and Proficiency Pay Programs," April 20, 1983.

[6] Everett, H.C., "A Lagrangian Algorithm for the Multiple Choice Integer Program," *Operations Research*, **11**, 299–317 (1963).

[7] Fisher, M.L., "Lagrangian Relaxation Method for Solving Integer Programming Problems," *Management Science*, **27**, 1–18 (1981).

[8] Garfinkel, R.S., and Nemhauser, G.L., *Integer Programming*, Wiley, New York, 1972, pp. 108–152.

[9] General Research Corporation, "ADP Support Services for Enlisted Systems," Contract Number MDA 903-87-C-0797, 15 March 1988.

[10] Geoffrion, A.M., "Lagrangian Relaxation for Integer Programming," *Mathematical Programming Study*, **2**, 82–114 (1974).

[11] Lovell, C.A.K., and Morey, R.C., "The Allocation of Consumer Incentives to Meet Simultaneous Sales Quotes: An Application to U.S. Army Recruiting," *Management Science*, **37**, 350–367 (1991).

[12] Streff, M.J., "The Estimation of United States Army Reenlistment Rates," M.S. thesis, Naval Postgraduate School, September 1989.
[13] Sweeney, D.J., and Murphy, R.A., "Branch and Bound Methods for Multi-Item Scheduling," *Operations Research*, **29**, 853–864 (1981).